On the computation of Mordell-Weil and 2-Selmer Groups of Elliptic Curves

J. E. Cremona

1 Introduction

Let E be an elliptic curve defined over $\mathbb{Q}$. In this note we present related methods to do the following tasks:

1. Prove that a given finite set of points in the Mordell-Weil group $E(\mathbb{Q})$ is independent;

2. Make the group law in the 2-Selmer group $S^2(E/\mathbb{Q})$ explicit, and hence show that a given finite set of elements in $S^2(E/\mathbb{Q})$ is independent.

The first provides an alternative to computing the height pairing matrix of the given set of points and showing that its determinant is non-zero. While that is easily done, for curves of large rank it requires some delicate consideration of precision in order to be sure of the result. The method here, by contrast, involves only "discrete" computations: finding roots of cubics and evaluating quadratic characters modulo primes. It was also described by Silverman in [6], attributed there to Brumer and myself. In fact, Brumer described the method to me in 1996; it was apparently used by him and Kramer in verifying the examples in [2], though the method is not explicitly mentioned there; so the method goes back to 1975 at least. We give it here as it is closely related to, and leads to, our second section where we apply similar ideas to 2-Selmer groups. We illustrate the method with the Martin-McMillen curve which has 23 independent points.

The second problem arises when doing explicit 2-descents on elliptic curves with no 2-torsion, as implemented in our program `mwrank`. Following the method set out in [1], we represent elements of the Selmer group $S^2(E/\mathbb{Q})$ by quartics $g(X) \in \mathbb{Z}[X]$ such that the genus 1 curve $Y^2 = g(X)$ is a 2-covering of E. These quartics are found by a finite search procedure. In [1], the resulting set of (equivalence classes of) quartics is treated as a set, without making explicit its structure as an elementary abelian 2-group. Indeed, one check on the calculations is to make sure that the size of the set obtained is a power of 2. We will show how to make explicit use of the group structure on $S^2(E/\mathbb{Q})$, via a homomorphism to $(\mathbb{Z}/2\mathbb{Z})^M$ for some $M > 0$. This has a number of practical advantages in terms of the running time of the resulting algorithm: we do not need to check equivalences between the quartics found; for a curve of rank r, we only find and consider r quartics instead of 2^r, saving much time in the search for rational points on the associated 2-coverings; and the search itself is made faster, since for every quartic we find, the remaining part of the search region is reduced by a factor of 2.

Our solutions to both problems generalize immediately to elliptic curves defined over general number fields; this is particularly true of the first, where it is likely to be considerably simpler than implementing the height pairing computation. We restrict to $\mathbb{Q}$ for ease of exposition.

2 Mordell-Weil Groups

The methods rely on explicitly embedding the groups $E(\mathbb{Q})/2E(\mathbb{Q})$ and $S^2(E/\mathbb{Q})$ into a direct sum $\oplus_{p\in S}E(\mathbb{Q}_p)/2E(\mathbb{Q}_p)$, where S is a finite set of "good" primes. Let E be given by an equation

$$E: \qquad Y^2 = f(X) = X^3 + AX^2 + BX + C,$$

where $f(X) \in \mathbb{Z}[X]$ has discriminant $\Delta \neq 0$. We may assume that this model is minimal at all odd primes, and we define a "good" prime to be one not dividing 6Δ. The Galois group G of f is a subgroup of S_3, and even of A_3 if Δ is a perfect square. We have $E(\mathbb{Q})/2E(\mathbb{Q}) \cong (\mathbb{Z}/2\mathbb{Z})^{r+t}$, where r is the rank of $E(\mathbb{Q})$ and $t = 0$, 1 or 2 according as $f(X)$ has 0, 1 or 3 rational roots.

Let p be a good prime. Since $p > 2$ and E has good reduction at p, we have

$$E(\mathbb{Q}_p)/2E(\mathbb{Q}_p) \cong E(\mathbb{F}_p)/2E(\mathbb{F}_p) \cong (\mathbb{Z}/2\mathbb{Z})^{k_p},$$

where $k_p = 0, 1$ or 2 according as the number of roots of $f(X)$ modulo p is 0, 1 or 3. If $f(X)$ is irreducible, then by the Chebotarev density theorem, the density of the primes with these behaviours is $\frac{1}{3}$, $\frac{1}{2}$, $\frac{1}{6}$ when $G = S_3$, and $\frac{2}{3}$, 0, $\frac{1}{3}$ when $G = A_3$. Since $E(\mathbb{Q}_p)/2E(\mathbb{Q}_p) \cong E(\mathbb{F}_p)/2E(\mathbb{F}_p)$ for good primes, in the rest of this section we use $\mathbb{F}_p$ in place of $\mathbb{Q}_p$, and our homomorphism can be defined via arithmetic modulo p.

Primes p for which $k_p = 0$ will be of no use to us since then the map $E(\mathbb{Q})/2E(\mathbb{Q}) \to E(\mathbb{F}_p)/2E(\mathbb{F}_p) = 0$ is trivial. It would be simpler to implement the method using only primes for which $k_p = 1$, but such an implementation would fail for curves with square discriminant (when there are no such primes), so we keep both types $k_p = 1$ and $k_p = 2$ in consideration.

It would also be possible, with more work, to use the local information at the primes 2 and 3, and also those of bad reduction. The results of [2] would be useful in those cases. However, for our practical purposes this is not necessary.

2.1 Definition of the maps ε_p and ε on $E(\mathbb{Q})/2E(\mathbb{Q})$

Let χ be the quadratic character modulo p, and $\psi : \mathbb{F}_p^* \to \mathbb{Z}/2\mathbb{Z}$ its additive version, so that $\chi(x) = (-1)^{\psi(x)}$ for $x \neq 0$.

First let p be a good prime for which $k_p = 1$, and let θ_p be the unique root of $f(X)$ modulo p. Then we define a map

$$\overline{\varepsilon}_p : E(\mathbb{F}_p) \to \mathbb{F}_p^* \to \mathbb{Z}/2\mathbb{Z}$$

by

$$\begin{aligned} P = (x, y) &\mapsto (x - \theta_p) \mapsto \psi(x - \theta_p); \\ P = (\theta_p, 0) &\mapsto f'(\theta_p) \mapsto \psi(f'(\theta_p)); \\ 0 &\mapsto 0. \end{aligned}$$

This map is a surjective group homomorphism with kernel precisely $2E(\mathbb{F}_p)$, and so gives an isomorphism

$$E(\mathbb{F}_p)/2E(\mathbb{F}_p) \cong \mathbb{Z}/2\mathbb{Z}.$$

Composing with the reduction homomorphism $E(\mathbb{Q}) \to E(\mathbb{F}_p)$ gives the desired map $\varepsilon_p : E(\mathbb{Q})/2E(\mathbb{Q}) \to \mathbb{Z}/2\mathbb{Z}$:

$$P \mapsto \varepsilon_p(P) = \overline{\varepsilon}_p(\overline{P}).$$

Writing $P \in E(\mathbb{Q}) \setminus \{0\}$ as $P = (u/w^2, v/w^3)$ with $u, v, w \in \mathbb{Z}$ and $\gcd(u, w) = \gcd(v, w) = 1$, set

$$\alpha(P) = \begin{cases} u - \theta_p w^2, & \text{if } u \not\equiv \theta_p w^2 \pmod{p}; \\ f'(\theta_p), & \text{if } u \equiv \theta_p w^2 \pmod{p}; \end{cases}$$

note that we cannot have $\alpha \equiv 0 \pmod{p}$, else θ_p is a double root of $f(X)$ modulo p, which cannot happen for p good. Then ε_p is given for $P \neq 0$ by

$$\varepsilon_p(P) = \psi(\alpha(P)) \qquad (P \neq 0).$$

Now let p be a good prime such that $k_p = 2$, and let $\theta_{i,p}$ for $i = 1, 2, 3$ be the (distinct) roots of $f(X) \pmod{p}$. For $P = (x, y) \in E(\mathbb{F}_p)$ with $2P \neq 0$, the elements $x - \theta_{i,p} \in \mathbb{F}_p$ are nonzero and their product is a square, and we define

$$\overline{\varepsilon}_p : E(\mathbb{F}_p) \to (\mathbb{F}_p^*/(\mathbb{F}_p^*)^2)^3$$

by

$$P = (x, y) \mapsto ((x - \theta_{1,p}), (x - \theta_{2,p}), (x - \theta_{3,p})) \pmod{(\mathbb{F}_p^*)^2},$$

the image lying in the subgroup H of $(\mathbb{F}_p^*/(\mathbb{F}_p^*)^2)^3$ of order 4 consisting of elements whose product is 1. We may extend $\overline{\varepsilon}_p$ to the points $(\theta_{i,p}, 0)$ of order 2 in $E(\mathbb{F}_p)$ as follows: if $x = \theta_{i,p}$, which can happen for at most one value of i since p is good, replace the component $x - \theta_{i,p}$ by $f'(\theta_{i,p})$. Then $\overline{\varepsilon}_p(P)$ still lies in H. Finally, $\overline{\varepsilon}_p(0) = 1$.

Again, $\overline{\varepsilon}_p$ is a group homomorphism with kernel $2E(\mathbb{F}_p)$ and image H; projecting onto the first two coordinates and switching to additive notation for H, we therefore have an isomorphism

$$E(\mathbb{F}_p)/2E(\mathbb{F}_p) \cong H \cong (\mathbb{Z}/2\mathbb{Z})^2.$$

As before, we now compose with the reduction map to get $\varepsilon_p : E(\mathbb{Q})/2E(\mathbb{Q}) \to (\mathbb{Z}/2\mathbb{Z})^2$. Given $P = (u/w^2, v/w^3) \in E(\mathbb{Q}) \setminus \{0\}$, for $i = 1, 2, 3$ set

$$\alpha_i(P) = \begin{cases} u - \theta_{i,p} w^2, & \text{if } u \not\equiv \theta_{i,p} w^2 \pmod{p}; \\ f'(\theta_{i,p}), & \text{if } u \equiv \theta_{i,p} w^2 \pmod{p}. \end{cases}$$

Then we have

$$\varepsilon_p(P) = (\psi(\alpha_1(P)), \psi(\alpha_2(P))) \in (\mathbb{Z}/2\mathbb{Z})^2.$$

Finally, let $S = \{p_1, p_2, \ldots, p_m\}$ be a set of good primes for which $k_p > 0$, and set $M = \sum_{i=1}^{m} k_{p_i}$. Then we obtain a homomorphism

$$\varepsilon : E(\mathbb{Q})/2E(\mathbb{Q}) \to (\mathbb{Z}/2\mathbb{Z})^M$$

by setting $\varepsilon(P) = (\varepsilon_{p_1}(P), \ldots, \varepsilon_{p_m}(P))$. This map is injective provided that m is large enough, from the following result.

Lemma 2.1. *Let $P \in E(\mathbb{Q}) \setminus 2E(\mathbb{Q})$. Then there exists a good prime p such that $\varepsilon_p(P) \neq 0$.*

Proof. We must find a good prime p such that the reduction $\overline{P} \notin 2E(\mathbb{F}_p)$. The X-coordinate of the points Q for which $[2]Q = P$ in $E(\overline{\mathbb{Q}})$ are the roots of a quartic polynomial $g(X) \in \mathbb{Q}[X]$. If $P \notin 2E(\mathbb{Q})$, then this quartic has no rational roots. Hence there exist (infinitely many) primes p such that $g(X)$ has no roots modulo p, and for such p we have $\overline{P} \notin 2E(\mathbb{F}_p)$ as required. □

2.2 Proving independence of points

Let $P_1, P_2, \ldots, P_n$ be rational points on the elliptic curve E which we wish to prove are independent. Take a finite set of good primes p_i and define the map $\varepsilon : E(\mathbb{Q}) \to (\mathbb{Z}/2\mathbb{Z})^M$ as before, factoring through $E(\mathbb{Q})/2E(\mathbb{Q})$. For $1 \leq i \leq n$ set $v_i = \varepsilon(P_i)$. Since ε is a homomorphism, if the P_i are dependent then this dependence relation will also be satisfied by the vectors v_i. So if we can show, using linear algebra over $\mathbb{F}_2$, that the v_i are linearly independent, then the points P_i are independent in $E(\mathbb{Q})/2E(\mathbb{Q})$, and hence also in $E(\mathbb{Q})$.

For this strategy to be successful in proving that a set of independent points is independent, two conditions must be satisfied. First, the number of primes m must be large enough for ε to be injective. Second, the points P_i must be independent in $E(\mathbb{Q})/2E(\mathbb{Q})$, which is a stronger condition than being independent in $E(\mathbb{Q})$. So if we find that a linear independence modulo 2 holds between the vectors v_i, we first increase the number of primes, so adding extra coordinates to the v_i; if the extended vectors are no longer linearly dependent, we will have succeeded in showing that the P_i are indeed independent. However, if the linear relation between the v_i persists when m is increased, it suggests that a linear relation holds between the P_i modulo $2E(\mathbb{Q})$.

In this case we will have an explicit linear combination $Q = \sum_{i=1}^{n} c_i P_i$ with each $c_i \in \{0, 1\}$, not all zero, such that $\varepsilon(Q) = 0$. We may determine whether $Q \in 2E(\mathbb{Q})$; if not, increasing m sufficiently will succeed in proving that the P_i are independent after all. If $Q = 0$, we have found a dependence relation between the P_i. On the other hand, if $Q = 2R \neq 0$, say, then we may replace one of the P_i for which $c_i = 1$ with R and repeat the process. After a finite number of steps we will succeed in either proving that the original points are independent, or find a relation between them.

See [6, Appendices D,G] for an alternative description of the algorithm, including a discussion of how many primes should be used and more details on how to find an explicit dependence relation between points which are not independent.

A program "`indep`" implementing this algorithm is available from the web site `http://www.maths.nottingham.ac.uk/personal/jec/ftp/progs/`.

2.3 Example

As an example of the method, we take the curve

$$\begin{aligned} E: \quad Y^2 + XY + Y = X^3 &- 19252966408674012828065964616418441723X \\ &+ 32685500727716376257923347071452044295907443056345614006, \end{aligned}$$

and the following 23 points in $E(\mathbb{Q})$:

$$
\begin{aligned}
P_1 &= (2509558762692426075, -4170888616355827764278386 28)
\end{aligned}
$$

This curve, and a different list of 23 independent points on it, were announced by R. Martin and W. McMillen in a posting by V. Miller to the `NMBRTHRY` email distribution list[1] in March 1998. The points we give here are not the same ones as originally posted, but they generate the same subgroup of $E(\mathbb{Q})$. We have LLL-reduced the height-pairing matrix of the original points to obtain this basis, whose elements have smaller heights (and considerably smaller denominators) than the original basis.

Using $m = 19$, with 19 primes between 7 and 157 giving $M = 23$, the dimension of the subspace of $(\mathbb{Z}/2\mathbb{Z})^{23}$ spanned by the images $\varepsilon(P_i)$ is only 22; adding a 20th prime 163 increases M to 24 and now the $\varepsilon(P_i)$ are independent in $(\mathbb{Z}/2\mathbb{Z})^{24}$, showing that the 23 points are independent in $E(\mathbb{Q})$. This computation takes around half a second on a 333MHz PC. If we compute the height pairing matrix to a precision of 38 decimals using Pari/GP, we find that its determinant is (approximately) $1.43 \cdot 10^{25}$; this computation takes about 2 seconds.

[1] See `http://listserv.nodak.edu/scripts/wa.exe?A2=ind9803&L=nmbrthry&F=&S=&P=970`

The table below gives the image under ε of the 23 points, using these 20 primes. The primes in bold face are those with $k_p = 2$, otherwise $k_p = 1$. The $\mathbb{F}_2$-rank of this matrix is 23.

p	**7**	31	43	47	**53**	59	67	71	83	89	97	109	113	127	**131**	139	149	**151**	157	163
P_1	1,0	1	1	1	1,1	1	0	0	0	1	0	0	0	0	1,1	0	0	0,1	1	1
P_2	0,1	1	0	0	1,0	1	1	1	0	1	1	1	0	0	0,1	0	0	0,1	1	1
P_3	1,1	1	0	0	1,1	0	1	1	1	1	1	1	0	0	0,0	0	0	0,1	1	0
P_4	1,0	0	1	1	1,1	1	1	0	1	1	1	0	0	1	0,0	0	1	1,0	1	0
P_5	0,0	1	1	0	1,1	0	1	1	0	1	0	1	0	0	1,1	0	0	1,0	1	0
P_6	0,0	1	1	0	1,0	0	0	0	1	1	1	1	0	1	1,1	0	0	1,1	0	1
P_7	1,0	0	0	1	1,1	1	1	1	0	0	1	0	0	1	0,0	1	1	0,0	0	1
P_8	1,1	0	0	1	0,0	0	1	1	1	1	1	1	1	1	0,1	1	0	1,0	1	0
P_9	0,1	0	1	0	0,1	1	0	0	1	0	0	1	1	1	0,0	1	1	1,0	1	1
P_{10}	1,0	1	0	1	1,1	1	1	1	0	1	0	1	1	1	1,1	0	1	1,0	0	1
P_{11}	1,0	0	1	0	1,1	1	0	1	0	1	0	0	1	1	0,1	1	0	0,0	1	0
P_{12}	1,0	0	0	1	0,1	1	1	0	1	0	0	0	0	1	0,1	0	1	0,0	0	1
P_{13}	0,1	1	0	1	0,1	1	0	0	0	1	1	0	0	1	1,1	0	0	0,0	1	1
P_{14}	0,1	1	0	0	1,1	1	1	0	0	0	0	1	0	1	0,1	1	1	1,1	0	0
P_{15}	1,0	0	1	0	1,1	1	1	1	0	0	0	1	1	0	1,0	1	1	1,1	0	1
P_{16}	1,1	0	0	1	0,1	0	0	0	0	0	0	0	0	1	0,1	0	1	0,1	1	0
P_{17}	0,1	0	0	0	1,1	1	0	0	1	0	1	0	0	1	0,1	0	0	0,1	1	0
P_{18}	1,0	0	1	1	1,1	1	1	0	0	0	0	1	1	1	1,0	1	0	1,0	1	1
P_{19}	1,0	0	0	0	0,0	1	0	0	0	1	1	1	1	0	1,1	1	1	0,0	1	0
P_{20}	1,0	0	1	0	0,1	0	1	1	0	1	0	1	1	1	0,1	1	1	1,1	1	0
P_{21}	1,1	0	0	0	0,1	0	0	1	0	1	1	1	0	0	0,0	0	0	1,0	0	0
P_{22}	1,0	0	1	0	1,1	0	1	1	1	0	1	1	0	1	1,1	0	0	1,1	0	1
P_{23}	0,0	0	0	1	1,0	1	1	1	0	1	0	1	1	0	1,0	1	1	0,1	1	0

3 Selmer Groups

The process of 2-descent on an elliptic curve E consists of embedding $E(\mathbb{Q})/2E(\mathbb{Q})$ into the 2-Selmer group $S^2(E/\mathbb{Q})$, which is also a finite elementary abelian 2-group. The cokernel of this embedding is the 2-torsion subgroup of the Tate-Shafarevich group.

The 2-descent algorithm has two phases: first, compute the Selmer group S^2; second, determine which elements in S^2 come from rational points. The first step is effective, but not the second, since there is (at present) no known way of deciding in all cases which elements of S^2 come from $E(\mathbb{Q})$ and which give non-trivial elements of Ш.

We give here a brief sketch of the algorithm; for more details, see [1] or [3].

Elements of S^2 may be represented by principal homogeneous spaces, curves of genus 1 which are 2-coverings of E, with affine equations of the form

$$\mathcal{C}_g : \quad Y^2 = g(X) = aX^4 + bX^3 + cX^2 + dX + e.$$

Here, $g(X)$ is a quartic polynomial, which we may take in $\mathbb{Z}[X]$, whose invariants I and J are related to the standard c_4 and c_6 invariants of a minimal model for E in a precise way (essentially, $I = c_4$ and $J = 2c_6$).

For a fixed pair (I, J) we search for all integer quartics $g(X)$ with these invariants, up to an equivalence which is defined so that two quartics are equivalent if and only if they determine the same element in S^2. Having found this (finite) set of quartics, we eliminate those for which the curve $\mathcal{C}_g$ does not have points in all completions of $\mathbb{Q}$, since these do not give elements of S^2, and then apply a test for equivalence to eliminate duplicates. This leaves us with a set of 2^s quartics g such that every element of S^2 is represented by exactly one homogeneous space $\mathcal{C}_g$, and we have determined the Selmer group, and its rank, s. Then we try to (somehow) determine which $\mathcal{C}_g$ have rational points and so come from rational points on E.

To search for the quartics, we first establish bounds on the leading coefficient a and the "seminvariant" $H = 8ac - 3b^2$: see [1] or [3], or [4] for improved bounds. Within the finite search region defined by these bounds, we look for pairs (a, H) satisfying the syzygy

$$H^3 - 48Ia^2H + 64Ja^3 = -27R^2, \tag{1}$$

which is satisfied by every quartic $g(X)$ with $R = b^3 + 8a^2d - 4abc$. In practice we use this syzygy as follows: for each of a finite set of auxiliary primes p, we precompute and store an array of 0-1 flags indexed by pairs $(a, H) \pmod p$, such that the flag is equal to 1 if and only if the left-hand side of (1) is -27 times a square modulo p. By restricting the search to those pairs (a, H) in the region for which all flags equal 1 we reduce the time for the search, since the only pairs we consider are very likely to lead to a solution to the syzygy and hence to a suitable quartic.

Under this scheme, we do not make any use of the group structure of S^2: the fact that the number of equivalence classes of locally soluble quartics turns out to be a power of 2 is not used, except as a check on the computation.

We now show how to define a homomorphism $\varepsilon = (\varepsilon_p)$ from $S^2(E/\mathbb{Q})$ to $(\mathbb{Z}/2\mathbb{Z})^M$ which extends the map defined in the previous section (for $p > 3$). That such an extension should exist is clear, since each $\mathcal{C}_g \in S^2$ has (by definition) a $\mathbb{Q}_p$-rational point for each prime p, and hence determines a well-defined class in $E(\mathbb{Q}_p)/2E(\mathbb{Q}_p)$. However, we will see that it is possible to define ε_p directly on the quartic $g(X)$, and even as a function of the pair (a, H) modulo p. The significance of this is that we can use ε during the syzygy sieving itself, to ensure that for each quartic $g(X)$ we construct, the vector $\varepsilon(g)$ is independent (over $\mathbb{F}_2$) of all previous ones. We do this by carefully adjusting the array of sieving flags as we proceed. The effect is that, for a curve of Selmer rank s, we only compute s quartics, which are automatically independent. This means that we no longer need to carry out equivalence testing between quartics, and also that the number of quartic curves $\mathcal{C}_g$ on which we must search for rational points is reduced from 2^s to s, which is a considerable saving for curves of large rank.

3.1 Definition of the maps ε_p and ε on $S^2(E/\mathbb{Q})$

Let $g(X) = aX^4 + bX^3 + cX^2 + dX + e$ be an integer quartic with invariants I and J such that the curve $\mathcal{C}_g$ defines an element in the Selmer group $S^2(E/\mathbb{Q})$. Set $H = 8ac - 3b^2$. The syzygy (1) may be written

$$F(H, 4a) = (H - 4a\varphi_1)(H - 4a\varphi_3)(H - 4a\varphi_3) = -27R^2$$

where $F(X,Y) = X^3 - 3IXY^2 + JY^3$ and φ_i for $i = 1,2,3$ are the roots of the (resolvent) cubic $F(X) = F(X,1) = X^3 - 3IX + J = 0$. In terms of the cubic $f(X) = X^3 + AX^2 + BX + C$ defining the curve E, we have $F(-4(3X+A)) = -2^6 3^3 f(X)$, so the roots θ_i of $F(X)$ (as defined in the previous section) are related to the φ_i by $\varphi_i = -4(3\theta_i + A)$, independently of the specific quartic $g(X)$.

First let p be a prime with $k_p = 1$, so that exactly one root θ exists modulo p, and correspondingly one value of φ. Set

$$\alpha(a,H) = \begin{cases} -3(H - 4a\varphi), & \text{if } H \not\equiv 4a\varphi \pmod{p}; \\ 3(H^2 - 16a^2 I) = F'(H,4a), & \text{if } H \equiv 4a\varphi \pmod{p}. \end{cases}$$

Then $\varepsilon_p : S^2(E/\mathbb{Q}) \to \mathbb{Z}/2\mathbb{Z}$ is given by

$$\varepsilon_p(g) = \varepsilon_p(a,H) = \psi(\alpha(a,H)).$$

That this is a well-defined homomorphism follows from [5, Proposition 3.2], if we note that our α is $9z$ in the notation of [5].

Now let p be a prime with $k_p = 2$, so that all three roots θ_i exist modulo p, with three corresponding values φ_i. Define $\alpha_i(a,H)$ for $i = 1,2,3$ as above, using φ_i in place of φ. Note that $\alpha_1\alpha_2\alpha_3$ is always a square; when $H \not\equiv 4a\varphi_i$ for all i, this product is non-zero and equal to $3^6 R^2$. Now we set

$$\varepsilon_p(g) = \varepsilon_p(a,H) = (\psi(\alpha_1), \psi(\alpha_2)) \in (\mathbb{Z}/2\mathbb{Z})^2.$$

For every good prime $p > 3$ we have now defined a homomorphism $\varepsilon_p : S^2(E/\mathbb{Q}) \to (\mathbb{Z}/2\mathbb{Z})^{k_p}$, in terms of the seminvariant pair (a,H) associated to a quartic $g(X)$ for which the curve $\mathcal{C}_g$ is a 2-covering of E representing an element of S^2. Putting these together for m primes p_i we thus obtain a map

$$\varepsilon : \quad S^2(E/\mathbb{Q}) \to (\mathbb{Z}/2\mathbb{Z})^M$$

where $M = \sum_{i=1}^m k_{p_i}$. For m (and hence M) large enough, this will be injective, for the same reasons as before (Lemma 2.1): a quartic represents the trivial element of S^2 if and only if it has a rational root, and maps to 0 under ε_p if and only if it has a root modulo p, so if g is a nontrivial quartic there will exist a prime p for which $\varepsilon_p(g) \neq 0$.

3.2 Some remarks on implementation

At the start of the computation (for a given elliptic curve), a supply of good primes p with $k_p > 0$ is determined, together with the corresponding value(s) of θ_p, and for each, a $p \times p$ array, indexed by pairs (a,H) modulo p, encoding whether quartics modulo p exist with this pair (a,H), determined from whether the left-hand side of the syzygy (1) is -27 times a square; this array also encodes the values of $\varepsilon_p(a,H)$. During the search, a pair of integers (a,H) is only considered if all the flags for this pair, for all p, are nonzero.

First suppose that $k_p = 1$ for all the primes used. When a (locally soluble) quartic g is found, we note the first "pivotal" prime p for which $\varepsilon_p(g) = 1$, and set all the (a,H)-flags modulo p such that $\varepsilon_p(a,H) = 1$ to zero. Then all future quartics g' found will have $\varepsilon_p(g') = 0$. This reduces the number of nonzero flags, so one effect is to speed up the remaining part of the search; we also ensure that the quartics found subsequently are necessarily independent of those found already, since the vectors $\varepsilon(g)$ for the quartics found at any given point are independent (being in echelon form with respect to a certain ordering of the primes).

For primes p with $k_p = 2$, the procedure is slightly more complicated. The first time such a prime is used as a pivot, we do the same as before, setting all the (a, H) flags such that $\varepsilon_p(a, H) = \varepsilon_p(g)$ to zero. If the same prime p is used as a pivot again, then the second time we set all the remaining (a, H) flags such that $\varepsilon_p(a, H) \neq 0$ to zero. Then all future quartics g' found will have $\varepsilon_p(g') = 0$, and the successive quartics found will be independent since the image vectors $\varepsilon(g)$ will by construction be in echelon form.

If the number of auxiliary primes is not large enough, it may happen that there will be no suitable pivotal prime at some point, as we may find a new quartic g with $\varepsilon(g) = 0$ which cannot be used as a pivot. Such quartics must be stored separately, and checked for equivalences; at the end, they form a subgroup $\ker(\varepsilon)$ of $S^2(E/\mathbb{Q})$, of order 2^{s_0} say, where $s_0 \geq 0$ depends on the number of primes used in the definition of ε. If we let s_1 denote the number of pivotal quartics found, which is the $\mathbb{F}_2$-dimension of $\mathrm{im}(\varepsilon)$, then the final value of the Selmer rank s is $s = s_0 + s_1$. Ideally, one should choose the number of auxiliary primes to be large enough so that ε is injective and hence $s_0 = 0$, but in a general-purpose program this is not practical. If we were to use sufficiently many primes to ensure that $\varepsilon : S^2(E/\mathbb{Q}) \to (\mathbb{Z}/2\mathbb{Z})^M$ was injective for curves of rank up to (say) 23, then the overheads of the necessary pre-computations would make the program rather inefficient for curves of very small rank.

Another complication arises from the fact that when a quartic g is found which is locally soluble, and hence represents an element of the Selmer group, we may not be successful in finding a rational point on the associated homogeneous space $\mathcal{C}_g$. This may happen either because $\mathcal{C}_g(\mathbb{Q})$ is empty, or because rational points exist on $\mathcal{C}_g$, but our search was not extensive enough to find any. We may then wish not to use this quartic as a pivot, as then we will later find other quartics g' in the same coset of the part of $E(\mathbb{Q})/2E(\mathbb{Q})$ which we have so far determined as g, and we may have more success in finding a rational point on $\mathcal{C}_{g'}$. More generally, there is a danger that the specific basis for the 2-Selmer group which we find may not be optimal, in the sense that we cannot tell in advance which homogeneous spaces will have rational points of smallest height. It is possible, using the methods of [5], to construct quartics representing all elements of the Selmer group from the quartics which form a basis; but this is not a trivial task, and we do not do this in our implementation. There does not seem to be an easy way of predicting which quartics will have the smallest rational points.

These considerations make the book-keeping involved in a complete implementation (such as our own program `mwrank`) rather intricate, but we will not go into further details here.

3.3 Example

Consider the elliptic curve

$$E : \qquad Y^2 = X^3 - 9217X + 300985,$$

which has rank 7, with no 2-torsion in Ш$(E/\mathbb{Q})$ so that the Selmer rank is also 7. The discriminant $\Delta_E > 0$, so the real points $E(\mathbb{R})$ form two connected components; the non-identity component (or "egg") contains rational points such as $(5, 505)$, so the total rank is $r = 1 + r'$ where r' is the rank of the connected component $E^0(\mathbb{Q}) = E(\mathbb{Q}) \cap E^0(\mathbb{R})$. We find r' by 2-descent.

Without using the map ε, the quartic search finds 64 inequivalent quartics (as well as a further 24 which are equivalent to these), each having rational points, which shows that $r' = \log_2(64) = 6$. Around 1400 equivalence tests have to be carried out between quartics during this computation.

Now using $m = 0, 1, 2, 3, \ldots, 9$ auxiliary primes $p \geq 5$ to define the filtering map ε, we find that the 2-rank of $\mathrm{im}(\varepsilon)$ increases from 0 up to 6 (on the identity component), while the order of $\ker(\varepsilon)$ decreases from 64 down to 1. When $m = 9$ (and $M = 11$), no equivalence tests are needed at all, since ε is injective. In addition, we only need to find rational points on 7 quartic homogeneous spaces instead of 127 (though for this curve all such points are very easy to find).

m	$\#\ker(\varepsilon)$	$\mathrm{rk}(\mathrm{im}(\varepsilon))$	r'	r
0	64	0	6	7
1	32	1	6	7
2	16	2	6	7
3	8	3	6	7
4	4	4	6	7
5	4	4	6	7
6	4	4	6	7
7	4	4	6	7
8	2	5	6	7
9	1	6	6	7

Remark. Regarding $\mathbb{R}$ as the completion of $\mathbb{Q}$ at the "infinite prime" $p = \infty$ allows us to view this example differently. When $\Delta > 0$, set $k_\infty = 1$ (otherwise $k_\infty = 0$), and define $\varepsilon_\infty : \quad E(\mathbb{Q})/2E(\mathbb{Q}) \to E(\mathbb{R})/2E(\mathbb{R}) \cong \mathbb{Z}/2\mathbb{Z}$ in the obvious way. Now add one component, given by the value of ε_∞, to the map $\varepsilon : E(\mathbb{Q})/2E(\mathbb{Q}) \to (\mathbb{Z}/2\mathbb{Z})^{M+1}$. In this way we may treat the infinite prime in the same way as the others. On the level of the Selmer group, when $\Delta > 0$ the quartics g considered also have positive discriminant, so have either 0 or 4 real roots, and we set $\varepsilon(g) = 1$ if and only if the number of real roots is 4. These are called quartics of "Type I" in [1] and [3]. In practice, the effect is that after finding one soluble quartic of Type I, we may restrict the subsequent search to quartics of Type II (with $\varepsilon(g) = 0$). This trick was already mentioned in [3], and all we have done here is to extend it from $\mathbb{R}$ to the p-adic completions of $\mathbb{Q}$.

A similar trick for the prime $p = 2$ also has important practical use, and will be discussed in a separate paper, since some new ideas are required. There may also be some benefit from using primes of bad reduction in the sieving process, but we have not investigated this.

References

[1] B. J. Birch and H. P. F. Swinnerton-Dyer, *Notes on elliptic curves I*, J. Reine Angew. Math. **212** (1963), 7–25.

[2] A. Brumer and K. Kramer, *The rank of elliptic curves*, Duke Mathematical Journal **44** (1977), no. 4, 715–743.

[3] J. E. Cremona, *Algorithms for modular elliptic curves*, second ed., Cambridge University Press, 1997.

[4] ———, *Reduction of binary cubic and quartic forms*, LMS J. Comput. Math. **2** (1999), 62–92.

[5] ———, *Classical invariants and 2-descent on elliptic curves*, Journal of Symbolic Computation (2000), to appear.

[6] J. H. Silverman, *The xedni calculus and the elliptic curve discrete logarithm problem*, Design, Codes, and Cryptography **10** (2000), 5–40.